



Technology Overview

November 2003

David Johnson
johnsond@redhat.com
214.435.2256



Agenda


- Red Hat Enterprise Linux Product Family Overview
- Red Hat Enterprise Linux 3 Technical Enhancements
- Red Hat Applications
- Red Hat Network




Traditional Red Hat Linux Release Model

Developed to meet the needs of the Open Source movement and early technology adopters

- 4-6 month release cycle
- Based on snapshot of core tree
- Latest open source technology
- ABI/APIs may change
- 5 beta cycles over 10 weeks
- Schedule driven; features may slip
- Red Hat features/enhancements fed back to core open source tree
- Replicate freely
- Limited support



Red Hat Enterprise Linux Release Model



Designed for enterprise-class application deployment

- 12-18 month release cycle
- Extensive 6 month beta cycle
- ISV technology input
- Improved API/ABI stability for broad ISV application support
- Enterprise technology focus
- ISV/OEM/customer-driven schedule
- Red Hat enhancements and features fed back to core tree
- 5+ year version support



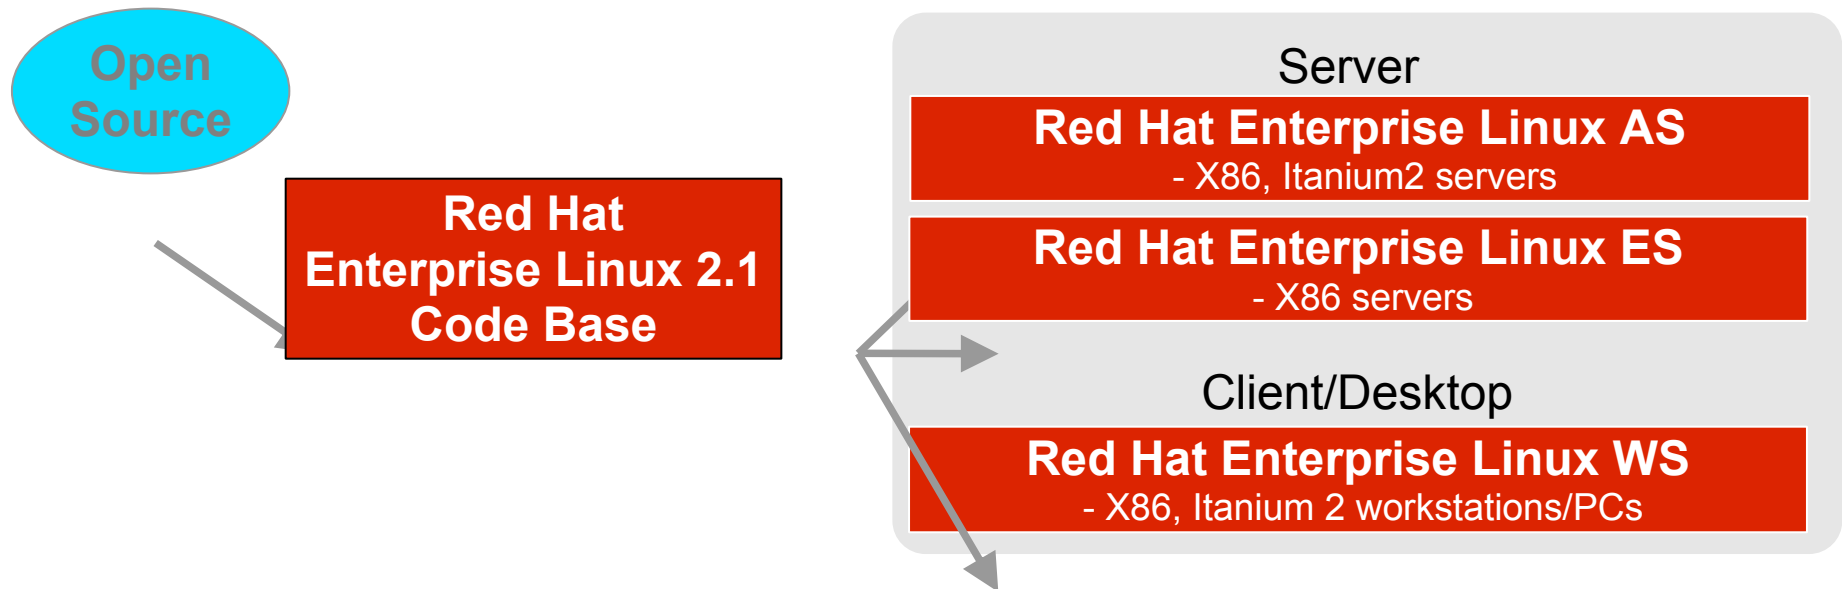
Fedora Project

Designed to be an incubator and proving ground for new technologies

- Open source project sponsored by Red Hat
- Used as a base for Red Hat supported products
- Provides rapidly evolving, technology-driven Linux distribution
- No explicit Security Errata/Support
- Critical fixes for 2-3 months after “release”
- Individual package developers QA and support
- No focus on binary compatibility



Enterprise Linux 2.1 Family – 2002-2003



**Stability & quality with
extended release cycle**



**Leadership price/performance
with audited benchmarks**



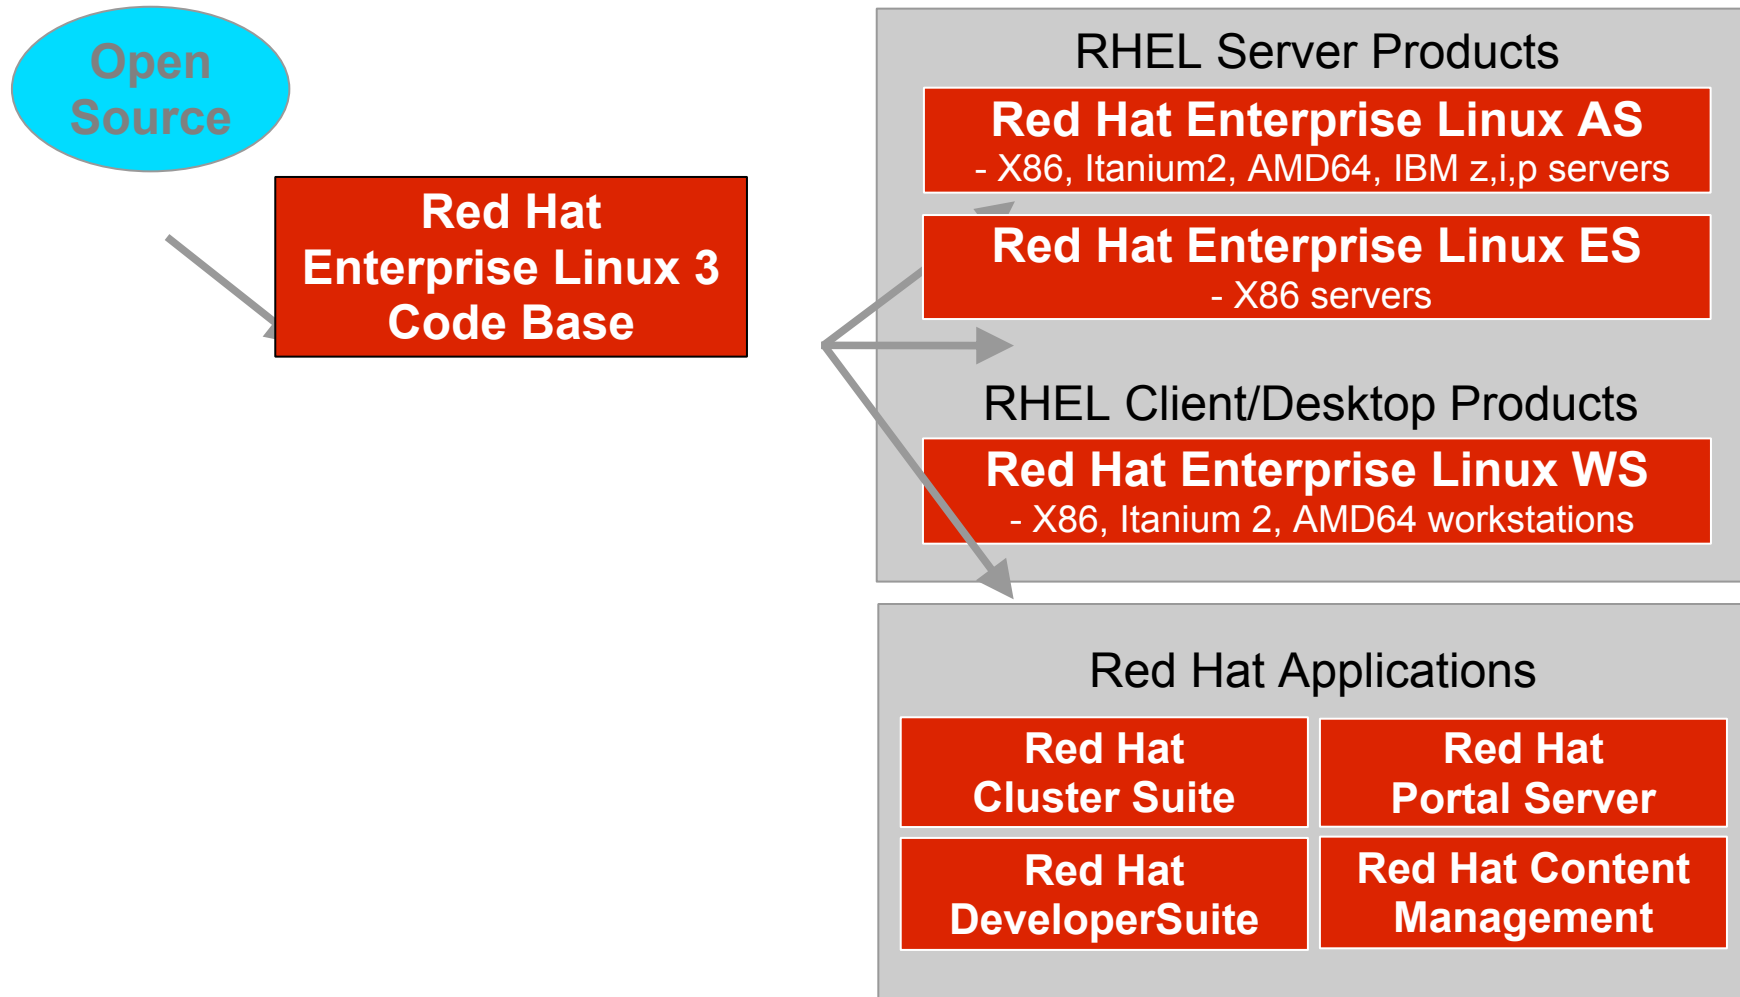
**Certified ISV applications
and OEM hardware**



**Services and Support from
Red Hat and Partners**



Enterprise Linux 3 Family – 2003-2004



Enterprise Linux Market Segmentation

| | Target Hardware | | Target Market |
|-----------------------------|---------------------|-------------------------------------|--|
| Red Hat Enterprise Linux AS | Advanced Server | >2 CPUs >8 GB memory (x86) | Large servers (e.g. database & enterprise applications) |
| Red Hat Enterprise Linux ES | Entry/Mid Server | 1-2 CPUs Up to 8 GB memory (x86) | SOHO & departmental servers (e.g. small-medium web, file & print configurations) |
| Red Hat Enterprise Linux WS | Workstation/Desktop | 1-2 CPUs | Corporate/home office productivity (e.g. Document creation, email, web, IM) Technical workstations (e.g. CAD/CAM, S/W development) HPC compute nodes |



Architecture Support Summary

Red Hat Enterprise Linux 3 provides broad architecture coverage:

| | Red Hat Enterprise Linux AS | Red Hat Enterprise Linux ES | Red Hat Enterprise Linux WS |
|---------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| Intel X86 compatible | Y | Y | Y |
| Intel Itanium | Y | N | Y |
| AMD AMD64 | Y | N | Y |
| IBM pSeries | Y | N | N |
| IBM iSeries | Y | N | N |
| IBM zSeries | Y | N | N |
| IBM S/390 | Y | N | N |



Red Hat Enterprise Linux 3 Technical Enhancements




Red Hat Enterprise Linux 3

The Red Hat Enterprise Linux 3 product family includes a large number of new features

- Over 100 Priority 1 features
- Over 350 general enhancements
- Requests from OEM and ISV partners, and customers
- Back-ported features from the Linux 2.5/2.6 kernel trees

A single source code base is used for all architectures

- Greatly improves code stability and maintainability
- 5 new architectures; 64-bit clean implementation
- Reduces feature skew; simplifies ISV application support



Enterprise Linux 3 Feature Summary

Focus on performance, scalability, availability, application development & standards support


- Kernel based on 2.4.21 with numerous 2.5/2.6 features
 - Better support for large SMP, memory and I/O configurations
- Forward compatibility between RHEL 2.1 and RHEL 3
- Greatly improved desktop environment
- 4GB-4GB Kernel/User Memory Split
- Enhanced standards support
- Enhanced security features
- Native Posix Threading Library
- GCC 3.2 tool chain environment
- Logical Volume Manager
- Diskless system support



4GB-4GB Split

Major new capability to support large physical memories and increased application virtual address space

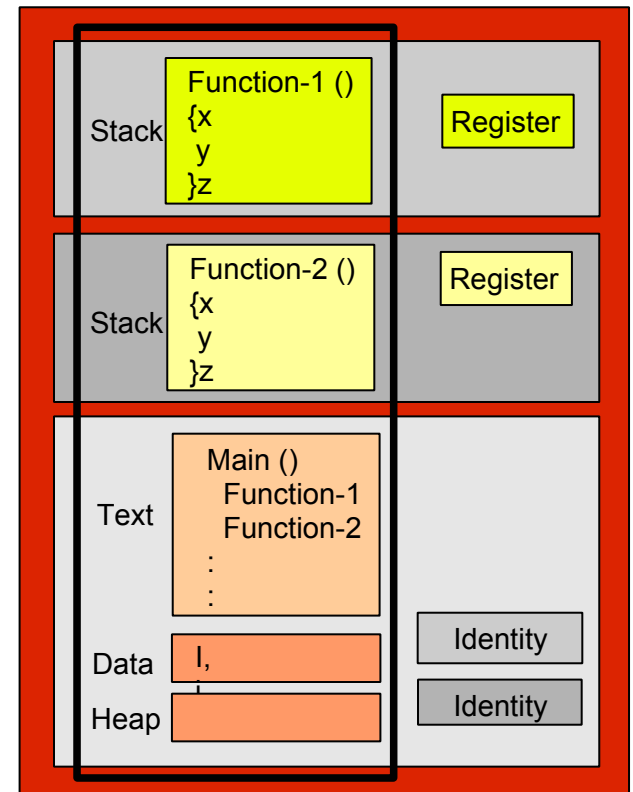
- A classic 32-bit 4GB virtual address space is split 3GB for user processes and 1GB for the kernel
- The new scheme permits 4GB of virtual address space for the kernel and almost 4GB for each user process



Native Posix Thread Library

Required for high performance multi-threaded commercial applications

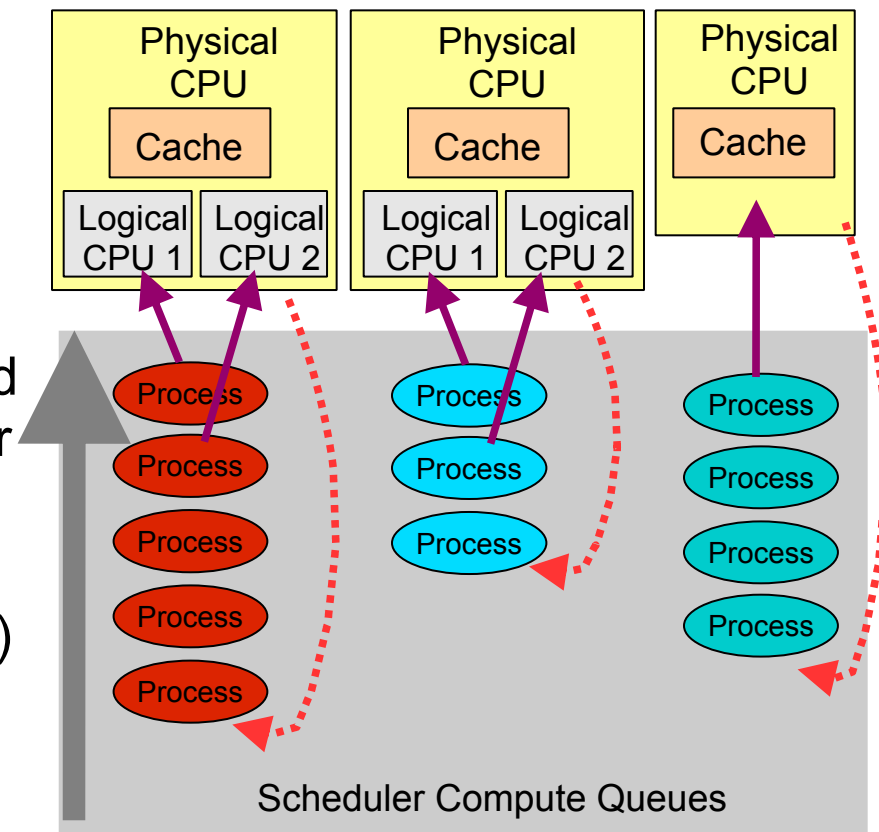
- e.g., Java
- Full implementation of POSIX threads
- Highly scalable, native implementation
 - Creation/deletion performance independent of the number of threads running
 - Informal benchmarks show >50,000 simultaneous thread creations-deletions per second
- Signal handling for multi-threaded processes
- Threaded core dumps



Hyperthreading Scheduler

Recognizes differences between logical and physical processors


- Optimizes process scheduling to take advantage of shared on-chip cache
- Implements one run queue per physical processor (as opposed to one run queue per processor or per system)
- Support for 16 logical CPUs (or 8 hyperthreaded CPU pairs)



Logical Volume Manager

Separate physical and logical devices


- ext2/ext3 filesystems resizable
- Allows flexible storage management
- Compatible with software RAID
- Uses LVM1 implementation (from Sistina)



Diskless System Support

Suitable for HPC, Telco and thin-client configurations

- Allows a Red Hat Enterprise Linux server to host other Red Hat Enterprise Linux images which clients boot over the LAN
- Minimal per-client storage overhead
- Clients can use local disks for swapping and general storage



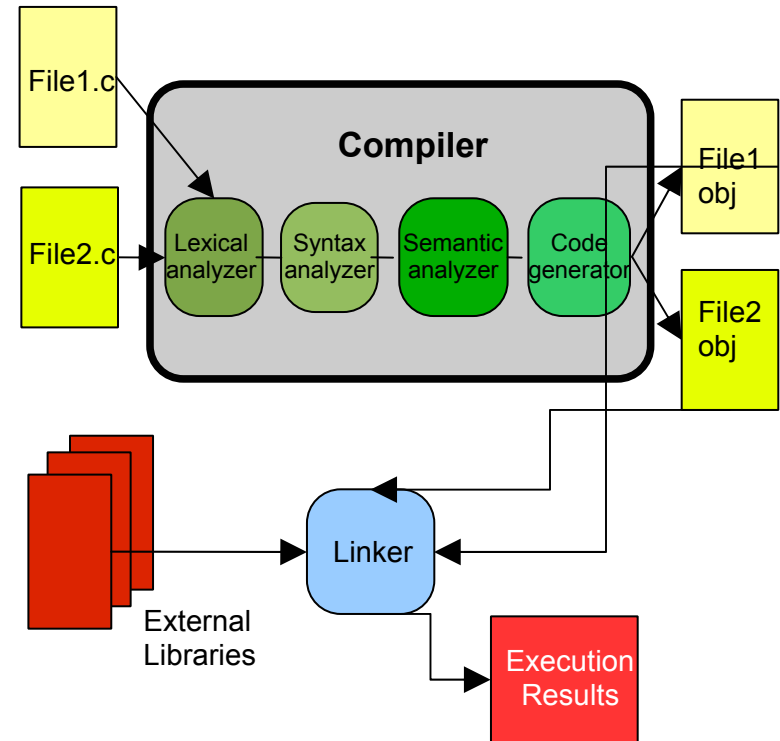
Compiler Environment

GCC 3.2 toolchain

- Full ANSI C++ support
- ISO C99 Standard support
- Memory debugging support
- Architecture Optimizations
 - Pentium IV s/w pipelining, etc
 - IA64 instruction scheduler
 - Compiler intrinsics for MMX & SSE (multimedia/streaming instructions)


GCC “ssa” toolchain included as a technology preview

- Static symbol assignment improves code generation



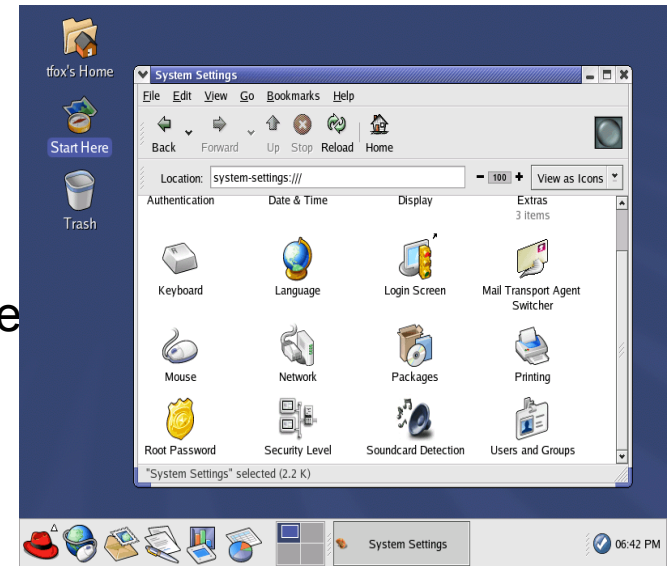
I/O Subsystem Improvements

- 64-bit SCSI/Fibre Channel DMA support (improved performance with >4GB memory)
- Up to 256 SCSI devices (permits larger systems to be configured)
- VaryIO support (permits larger I/O transfers)
- Serial ATA support - SATA1 (for Intel PIIX/ICH ATA)



Desktop Environment

- New graphics hardware support
 - XFree86 4.3.x
- BlueCurve graphical user interface
 - Unified GNOME/KDE look and feel
 - Designed for usability
- Bundled productivity applications
 - OpenOffice & Ximian Evolution
 - Mozilla 1.4 web browser



Serviceability

Kernel crash dump and analysis enhancements

- .Crash – core dump analysis suite
- .Configurable application core dump paths
- .Code profiling support included in the kernel – Oprofile
 - .System-wide profiler, capable of profiling all kernel/library/application code
 - .Uses hardware performance counters in the CPU
 - .Includes several post-profiling tools



Networking

- Improvements to channel bonding
 - Failover & bandwidth aggregation for servers with multiple NICs
- Kernel IPsec - secures IPv4 traffic
 - Tunnel mode builds tunnels between subnets
 - Transport mode secures communication directly between two machines
 - Packets are encrypted, authenticated & anti-replay protected
 - Able to communicate with IPsec devices and OS
- Kernel IPv6 support (more complete implementation than in 2.1)
- Kernel support for both IGMP V2 & V3 (Internet group management protocol)



NFS

Significantly improved stability & performance

- .Client-side focused performance enhancements
 - .NFSv3 readdirplus caches directory information
- .Enhanced robustness
- .NFS over TCP
- .O_Direct support added
- .ACL support with Red Hat Enterprise Linux 3



Security

Stronghold capabilities included in RHEL3 AS, ES & WS

- Previously provided as a separately purchasable product
- Updated to Apache 2.0 web server
- Includes OpenSSL, PHP, mod_perl, etc.



Java Support

- Native Java environment: GCJ / libgcj (Java GCC compiler front-end)
- Third-party Java environments included on 'Extras' CD and RHN channel:
 - IBM Java for x86, ia64, iSeries, pSeries, zSeries, S/390
 - Sun Java for x86 and ia64
 - BEA JRockIt for x86 and ia64



Security

File system ACLs

- .Unix file permissions not always adequate
 - .Multiple UIDs, Groups and set-UID apps proliferate.
- .ACLs are additional sets of read/write/execute triplets
- .Can be added to any objects
 - .Files, directories, devices, or any other file system objects
- .Highly configurable - fine-tune access
 - .Without resorting to multiple groups or set-UID apps
- .Includes support for NFS mounted file systems



Security

Kernel-level cryptography

- General purpose cryptographic api in the Kernel
 - Used by IPsec
 - Allows various types of cryptographic transforms to be used: Ciphers, digests, and compressors
 - Cryptographic algorithms are pluggable, e.g. AES, MD5
 - Can be implemented as loadable modules.
 - Not visible to user-level applications
- Potential support for Crypto-Hardware Possible
- Allows encryption to be done within the Kernel
 - Eliminates external patch to kernel



64-Bit Landscape

- AMD64 – 32-bit compatibility libraries + h/w execution
- IBM pSeries/iSeries – 32-bit compatibility libraries (ppc64)
- IBM zSeries/s390x
- Itanium2
 - Partial x86 compatibility libraries
 - Intel's execution layer not bundled



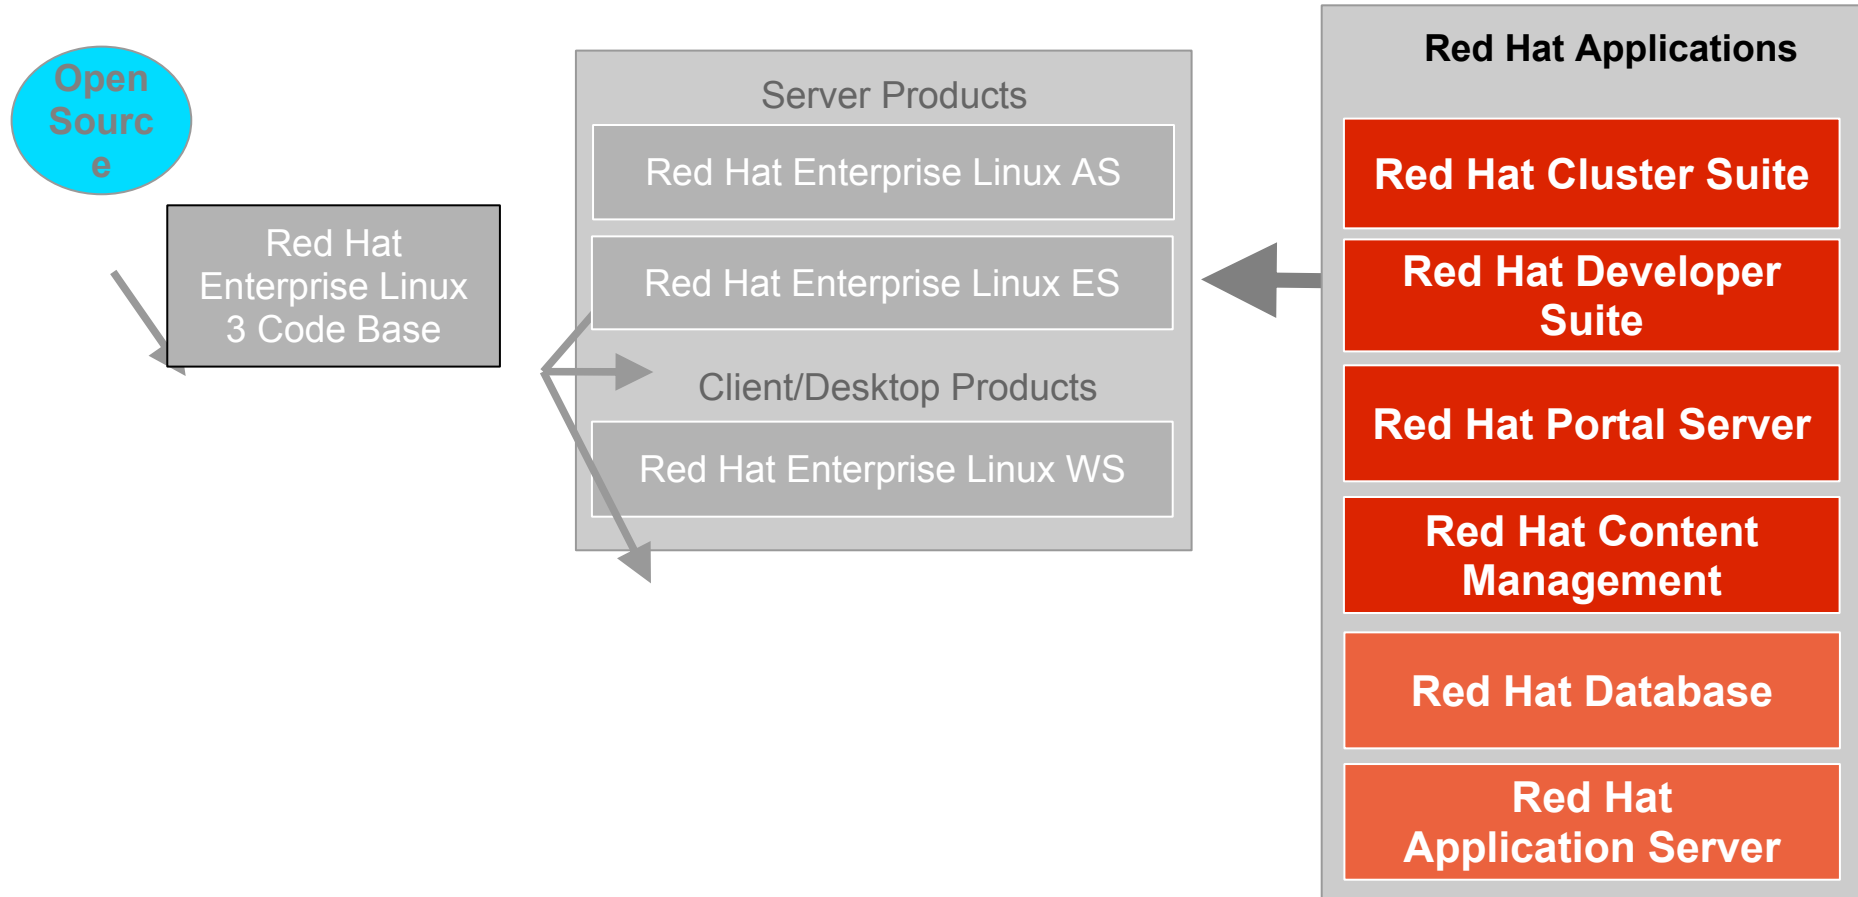
Standards Support

- Linux Standards Base 1.3 compliance (LSB)
 - Standard available at <http://www.linuxbase.org>
- NIAP Common Criteria certification expected to be complete by the end of 2003 (National Information Assurance Partnership)
 - Internationally accepted standard
 - Specified by US Department of Defense
 - <http://niap.nist.gov>



Red Hat Applications

Red Hat Applications



Red Hat Applications

Red Hat has extended the value of open source solutions by developing a suite of layered products for Enterprise Linux

- Cluster Suite - high availability “failover” clustering
 - Previously available only with Enterprise Linux AS
 - Now available as a layered product for AS & ES
 - Supported on x86-compatible systems
- Developer Suite
 - Eclipse-based IDE & developer tools
 - Supported on x86-compatible systems
- Other products are under active development for delivery in 2004



Categories of Clustering

Load Balancing clusters

- .Provide cost effective scalability of network-based services
- .Delivered by “IP Load Balancing” (aka Piranha) technology in Cluster Suite

High availability clusters

- .Provide increased availability of server & network-based applications
- .Delivered by “Cluster Manager” technology in Cluster Suite

High performance clusters (HPC)


- .Provide high performance for parallel applications
- .Not delivered by Cluster Suite



IP Load Balancing Clusters (Piranha)

Provides the capability for building scalable and available server farms

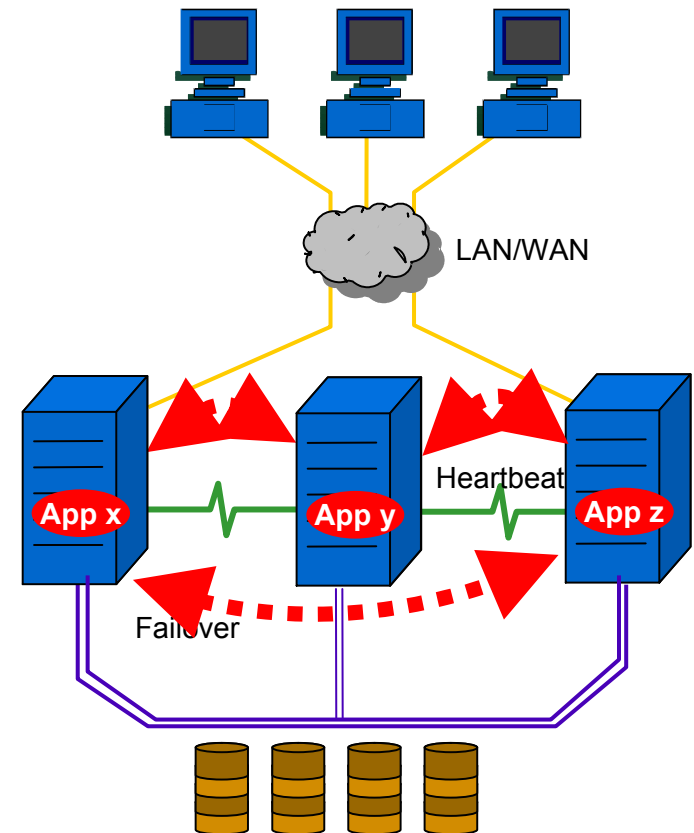
- Defense against capacity ceilings and service outages
- Load-balances incoming traffic
- Appears as a single virtual server
 - Transparent to clients
 - Insulated from failure
- Suitable for partitionable problems and applications (e.g., web servers)



High Availability Clusters

Increase the availability of non-partitionable, data-centric applications


- Database systems, file servers (NFS/CIFS), and mail servers
- Enhanced to support n-node failover clusters (up to 8 nodes)
- Applications run on any machine
- Shared SCSI or Fibre Channel data
- Allows scalability to large number of applications
- Configuration/management GUI



IP Load Balancing & HA Failover

Complementary clustering capabilities can be used together in a 3-tier architecture

- Virtual Server Clusters for scalable front end web processing
- Apache web servers
- HA Cluster for back end database and file serving



Cluster Management

Command line and GUI management interfaces

- Command line utilities for
 - Initial cluster configuration
 - Service configuration
 - Status checking
- Java console/applet for configuration and status monitoring

The screenshot displays the Cluster Management GUI for a cluster named 'Rainbow'. The status is 'Cluster is running' and the 'On Member' is 'red.lab.boston.redhat.com'. The 'Members' section lists six nodes: blue, cyan, green, magenta, red, and yellow, all with an 'Active' status. The 'Services' section shows a table of running services.


| Name | State | Member | Last Transition | Monitor Interval | Restarts |
|---------------|---------|---------|-----------------|------------------|----------|
| greenonly | Running | green | 12:50:17 Jul 28 | 0 | 0 |
| greenpref | Running | green | 12:50:17 Jul 28 | 0 | 0 |
| redonly | Running | red | 16:49:05 Jul 28 | 0 | 0 |
| redpref | Running | red | 16:49:05 Jul 28 | 0 | 0 |
| service_nfs_1 | Running | magenta | 16:56:33 Jul 28 | 0 | 0 |
| service_nfs_2 | Running | magenta | 16:57:39 Jul 28 | 0 | 0 |



Red Hat Developer Suite

Provides a complete development environment for Enterprise application developers

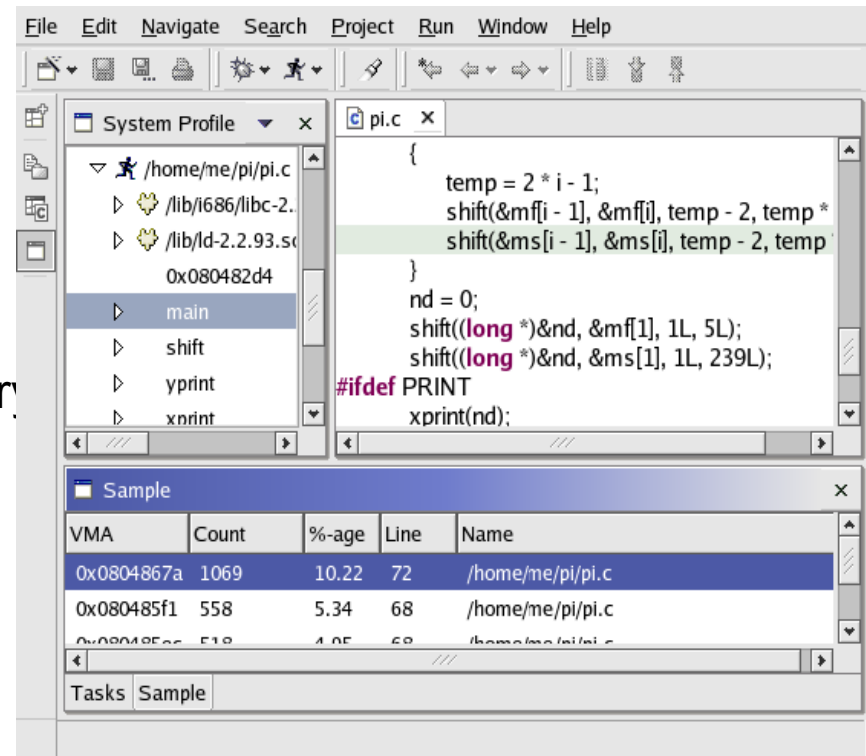
- Eclipse IDE framework & plugins
- Includes Eclipse plugins:
 - C/C++ Development Tools (CDT)
 - Java Development Tools (JDT)
 - RPM/SRPM
 - Oprofile
 - Plugin Development Environment (PDE)
- Supports a wide range of operating systems beyond Red Hat Enterprise Linux, including Windows (XP, 2000, 98, ME) and Sun Solaris 8



Red Hat Developer Suite

Compiled into a native binary using GCJ, the leading open source Java compiler

- Made possible by Red Hat's GCJ innovations and Eclipse enhancements
- Native compilation provides:
 - Freedom from proprietary JVMs
 - Faster startup
 - Faster operation overall



Web Applications

An Open Source web application framework based on the ArsDigita Community System

- .Red Hat Content Management System
- .Red Hat Portal Server
- .Red Hat Collaboration Applications



Red Hat Content Management System

Tools to create, manage and publish database-backed content

- Version Management – storage, archival, rollback and reuse of text, media, and arbitrary data
- Workflow – customizable role-based engine for tasking, routing, notification, and auditing of production process
- Lifecycle – automated publication/expiration management
- Notifications – tracking of task lists and email notifications based on workflow
- Syndication – standards-based content sharing capabilities
- Templates – allows abstract content to take on consistent look and feel
- Categorization – tool for logically grouping content



Visit www.redhat.com/software/rhea for more information



Red Hat Portal Server

Framework to create and administer personalized and group-oriented access to applications and information

- Integrated Entry Point – aggregates access ports to data and applications
- Personalization – Flexible portal layout editing, user-controlled application management
- Search and Navigation – Supports Intermedia and Lucene for full-text and binary search formats
- Application Development and Integration – Flexible Java API to develop custom applications and "portlet" views into existing systems
- Administration and Security – centralized administration interface with fine-grained user/group permissioning
- Collaboration – Integrates with Collaboration Suite to create shared workspaces and team room capabilities

Visit www.redhat.com/software/rhea for more information



Red Hat Collaboration Applications

Set of applications that enable people to work together in “context”

- .Enables asynchronous communication among distributed teams
- .Centralized shared file storage via Document Manager
- .Interactive tools for end-users including bulletin boards, chat, surveys, polls, etc.





Red Hat Network


Red Hat Network

Software Management Module

- Delivery of operating system and software applications
- Manage software updates
- Enable system configuration and profile management



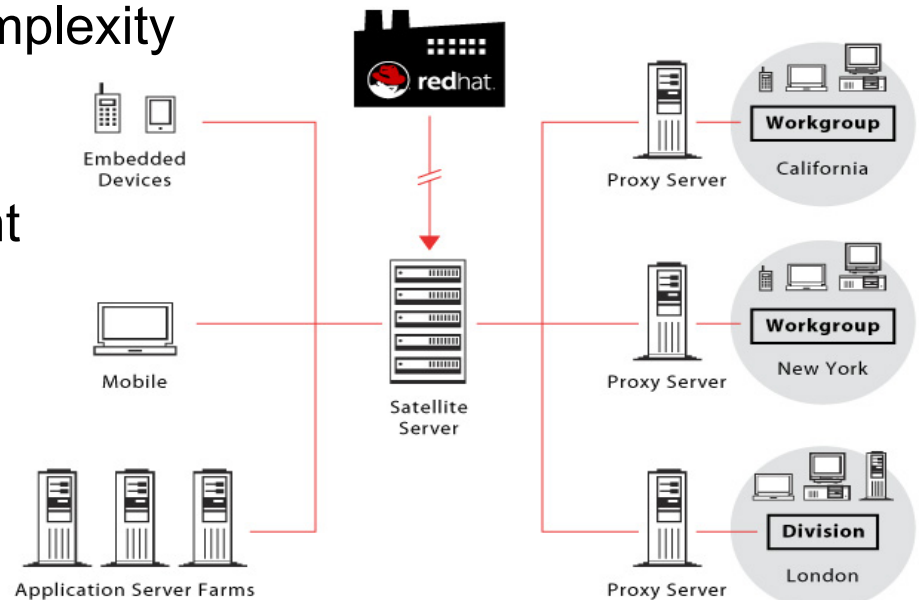
Management Backplane



RHN – Software Management

Delivery vehicle for maintenance to Red Hat Linux clients

- Improved system reliability and security
- Reduced management complexity
- Centralized control
- Cost-effective management of distributed systems
- On-demand access to Red Hat expertise



RHN – Software Management

How does it work?

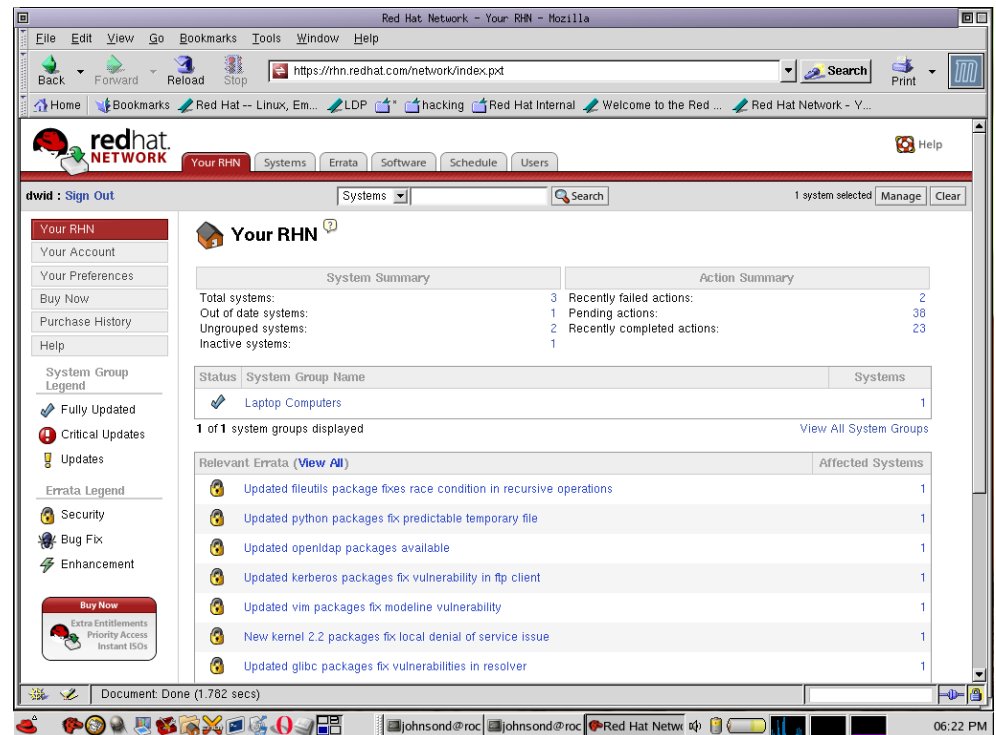
- Systems are registered with RHN
- System profiles stored in RHN database
- Entitlements provide “right” to receive service
- Actions scheduled by Administrator
- Systems check in via RHNSd
- Updates delivered as RPM packages



RHN – Software Management

Enterprise service demands enterprise features

- Group Management
- Channel Management
- User Management
- Activation keys
- Proxy
- Satellite



The screenshot displays the Red Hat Network (RHN) web interface within a Mozilla browser window. The address bar shows the URL <https://rhn.redhat.com/network/index.ppt>. The interface includes a navigation menu on the left with options like 'Your RHN', 'Systems', 'Errata', 'Software', 'Schedule', and 'Users'. The main content area is titled 'Your RHN' and features a 'System Summary' table, an 'Action Summary' table, and a 'Relevant Errata' section.

| System Summary | |
|----------------------|---|
| Total systems: | 3 |
| Out of date systems: | 1 |
| Ungrouped systems: | 2 |
| Inactive systems: | 1 |

| Action Summary | |
|-----------------------------|----|
| Recently failed actions: | 2 |
| Pending actions: | 38 |
| Recently completed actions: | 23 |

| Status | System Group Name | Systems |
|--------|-------------------|---------|
| ✓ | Laptop Computers | 1 |


1 of 1 system groups displayed [View All System Groups](#)

| Relevant Errata (View All) | Affected Systems |
|--|------------------|
| Updated fileutils package fixes race condition in recursive operations | 1 |
| Updated python packages fix predictable temporary file | 1 |
| Updated openldap packages available | 1 |
| Updated kerberos packages fix vulnerability in ftp client | 1 |
| Updated vim packages fix modeline vulnerability | 1 |
| New kernel 2.2 packages fix local denial of service issue | 1 |
| Updated glibc packages fix vulnerabilities in resolver | 1 |

Software Management Roadmap


Red Hat Network – Cactus

- Targeted for release in Q3 CY2003
- Transaction rollback processing (aka “undo”)
- Vastly improved scheduling for all tasks
- Finer-grained user roles




Hosted RHN

- Efficient consumption of bandwidth?
- System profiles stored at redhat.com



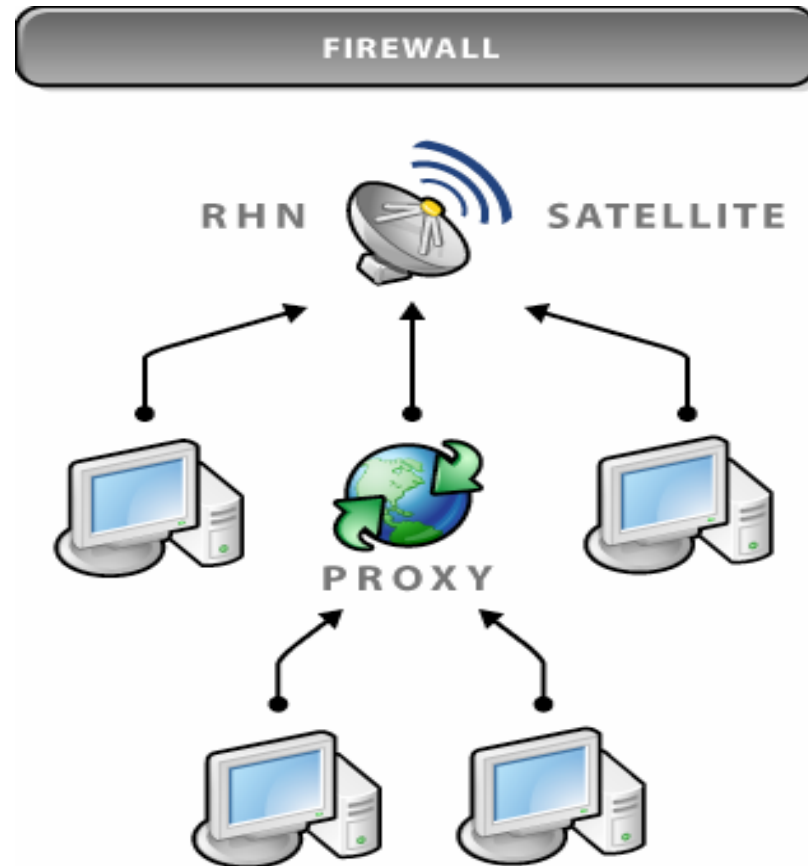
RHN Proxy Server

- Efficient consumption of bandwidth?
- System profiles stored at redhat.com

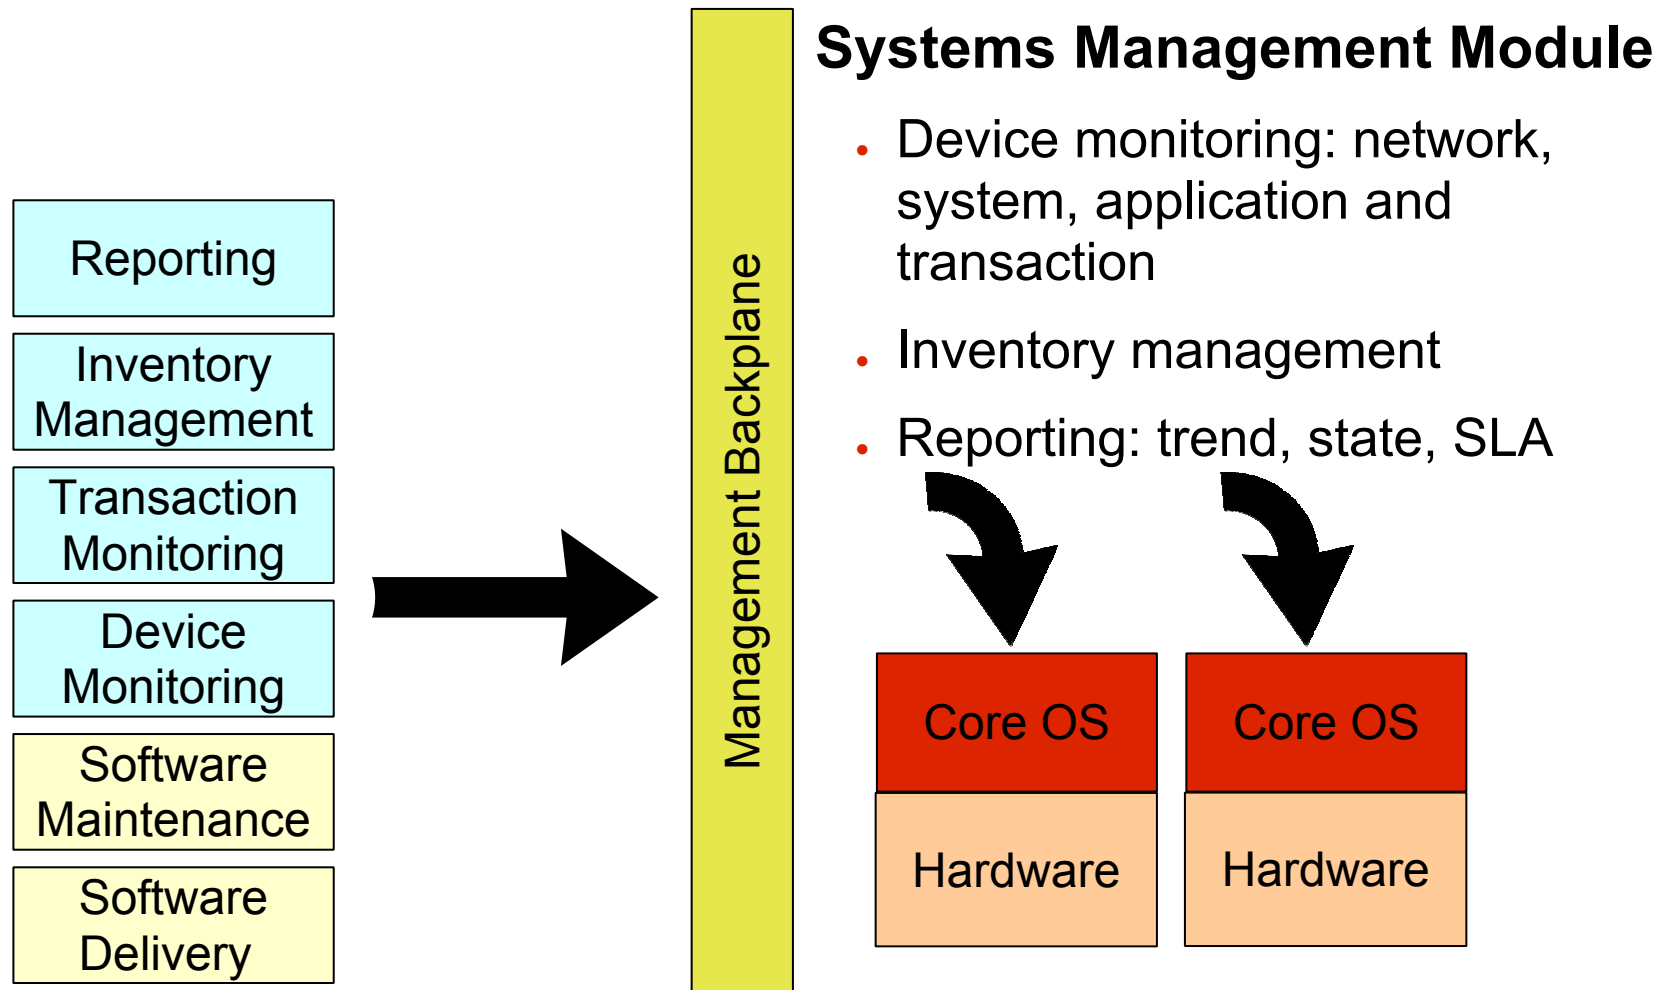


RHN Satellite Server

- Efficient consumption of bandwidth?
- System profiles stored within your organization



Red Hat Network



RHN – Systems Management

Monitoring Module

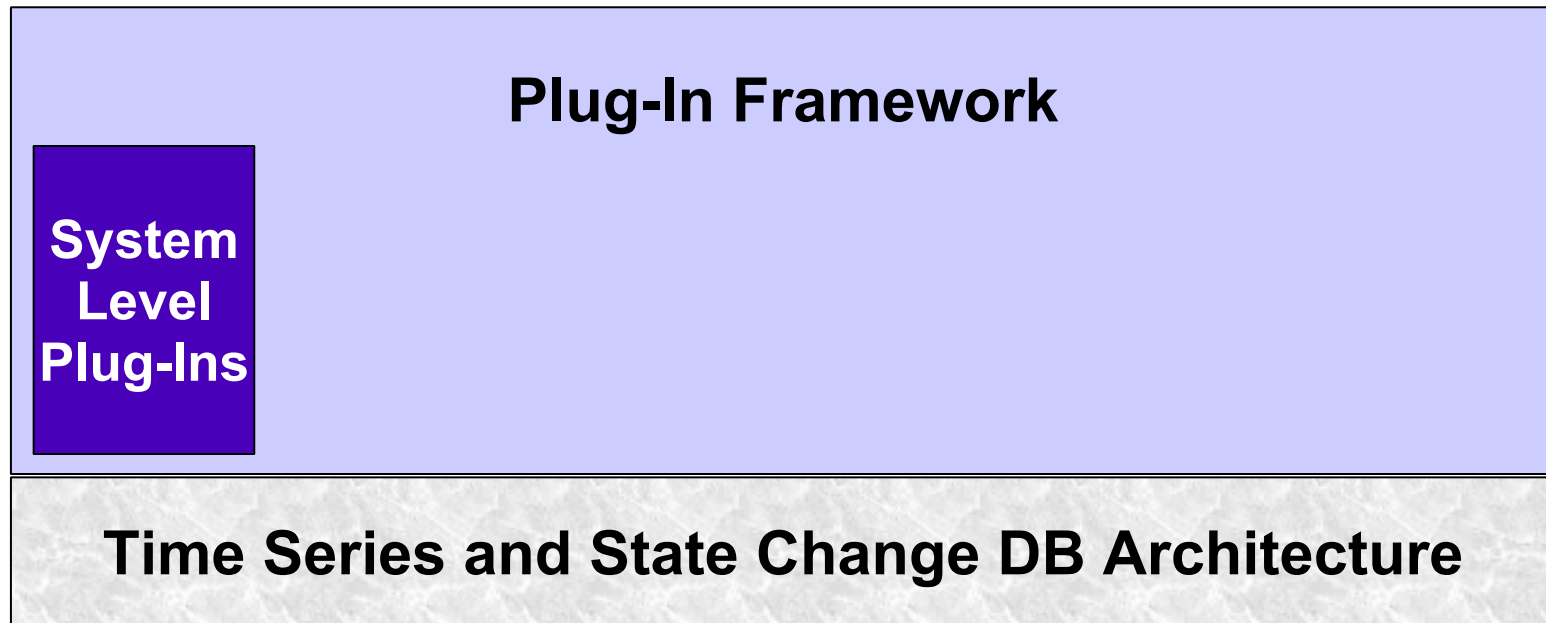
- Holistic view of components
- Web-based interface
- Support integration for MoM consoles
- Agent-less local software architecture
- Plug-in framework for fast support of device and application monitoring
- Multi-tenant backend
- Heterogenous OS monitoring



RHN – Plug-in Framework

System Level Plug-ins

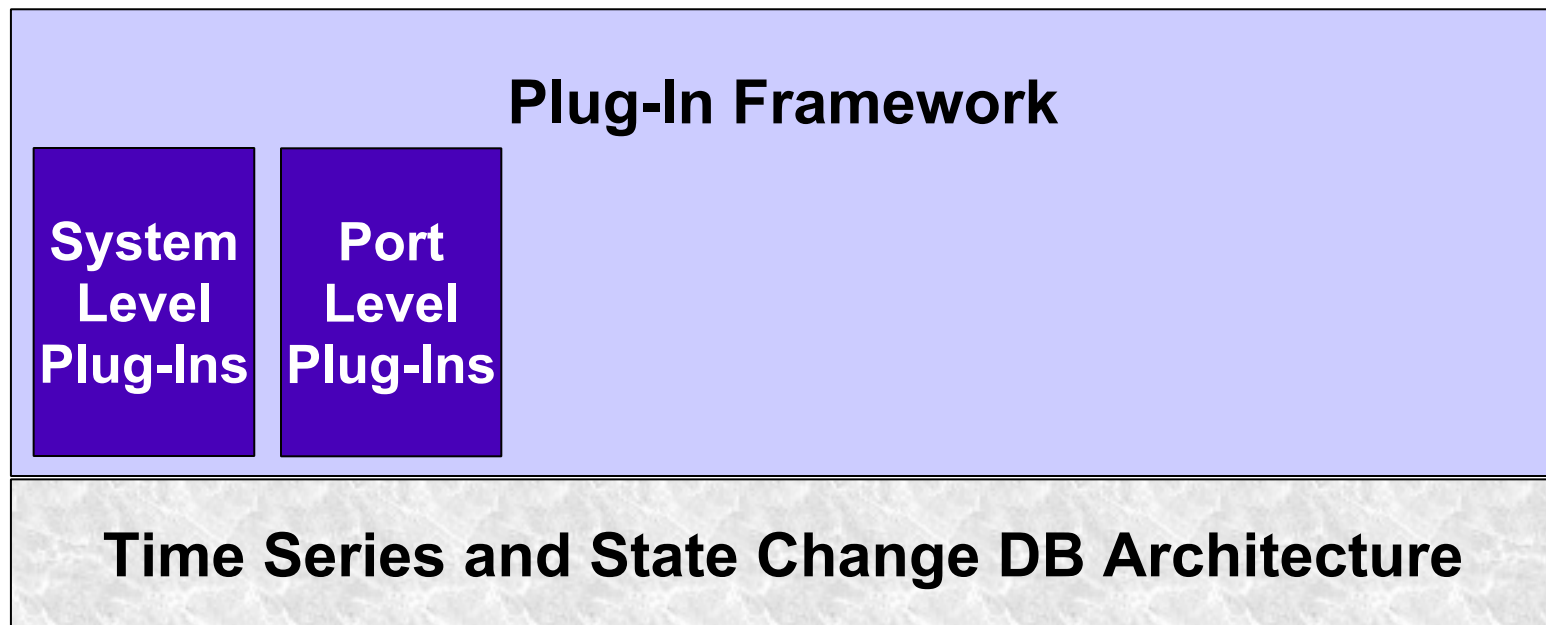
CPU, Disk, Ping, Memory, Load,
Users, Procs, Swap, etc.



RHN – Plug-in Framework

Port Level Plug-ins

HTTP, HTTPS, FTP, POP, IMAP,
SMTP, SNMP, IRCD, NNTP, DNS,
ORACLE, POSTGRES, REAL, etc.

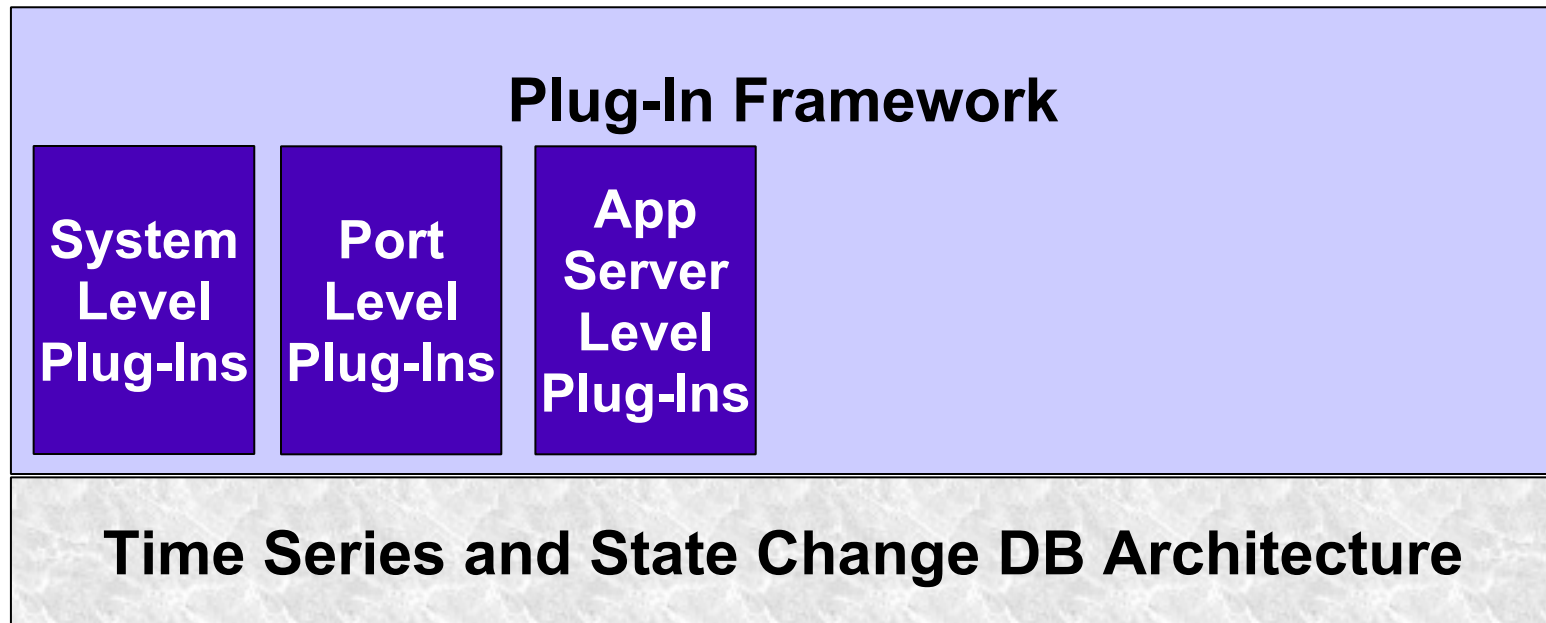


RHN – Plug-in Framework

Application Server Plug-ins

BEA Weblogic, IBM

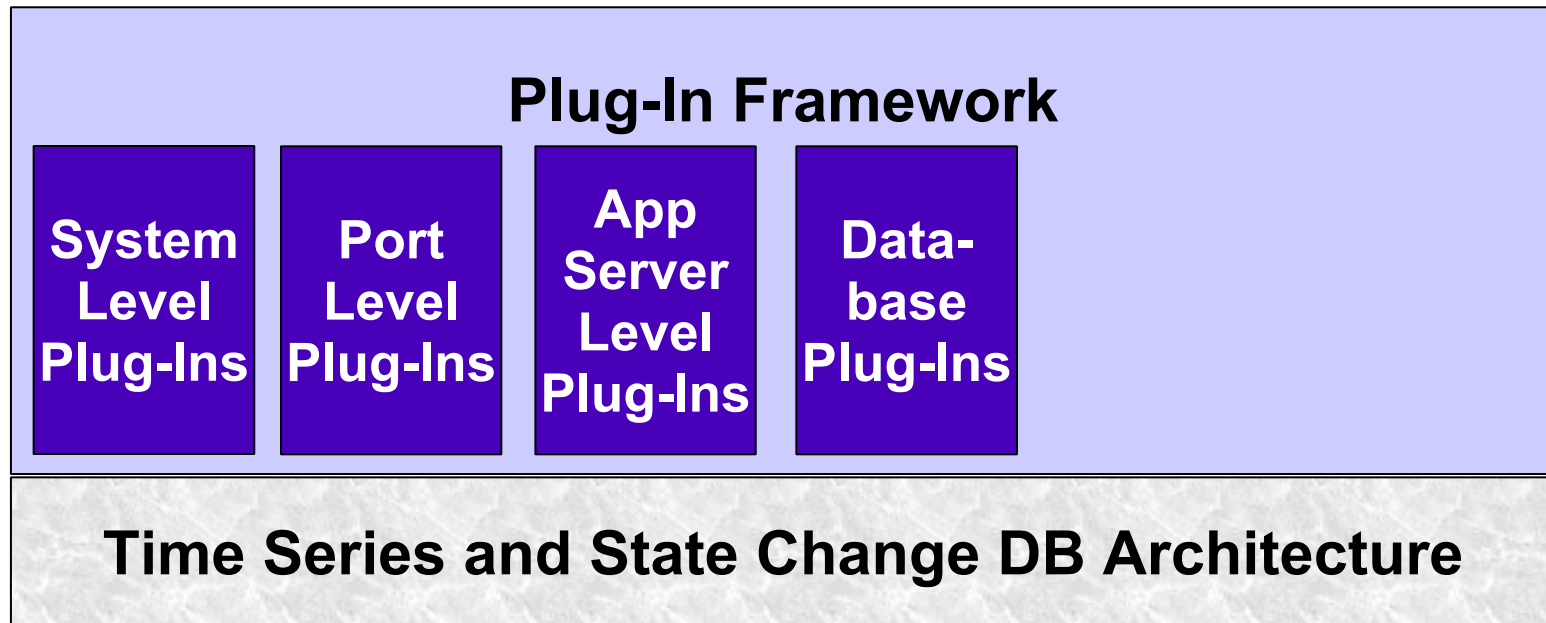
Websphere, ATG Dynamo, Sun ONE



RHN – Plug-in Framework

Database Plug-ins

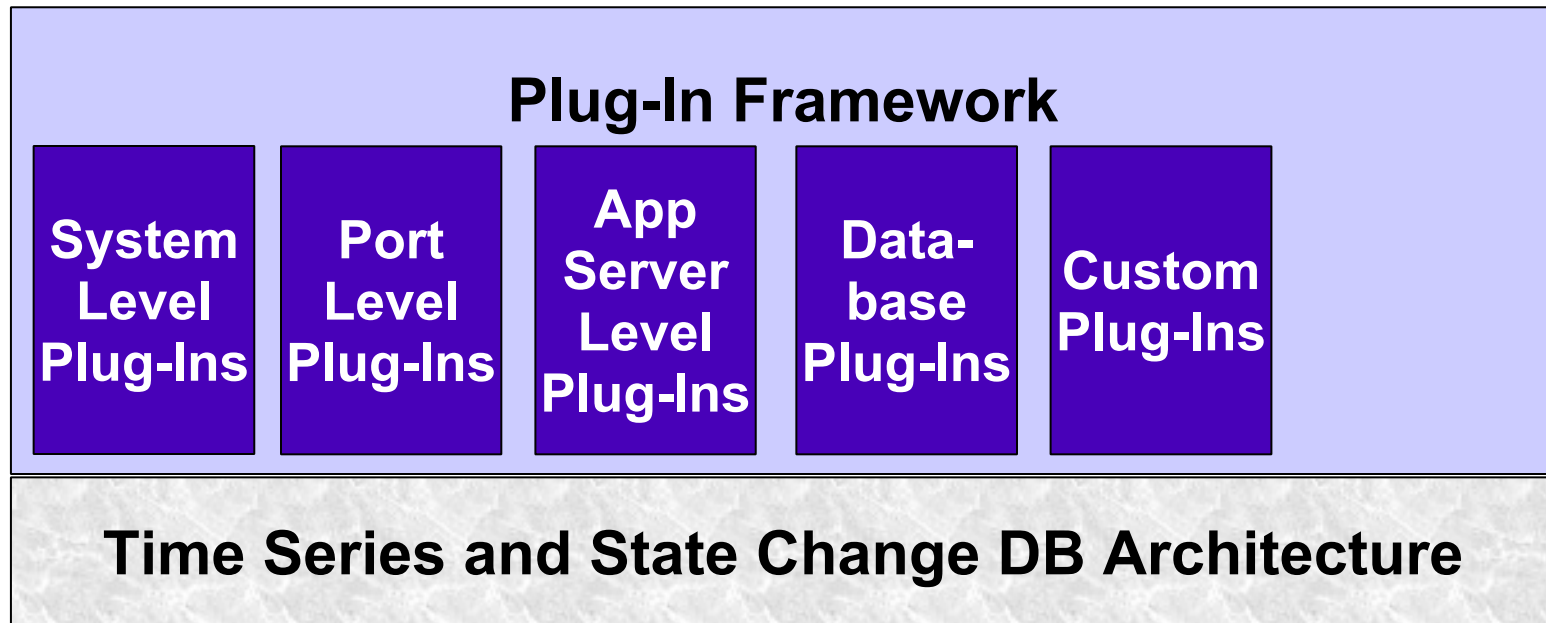
Oracle, SQLServer,
MySQL



RHN – Plug-in Framework

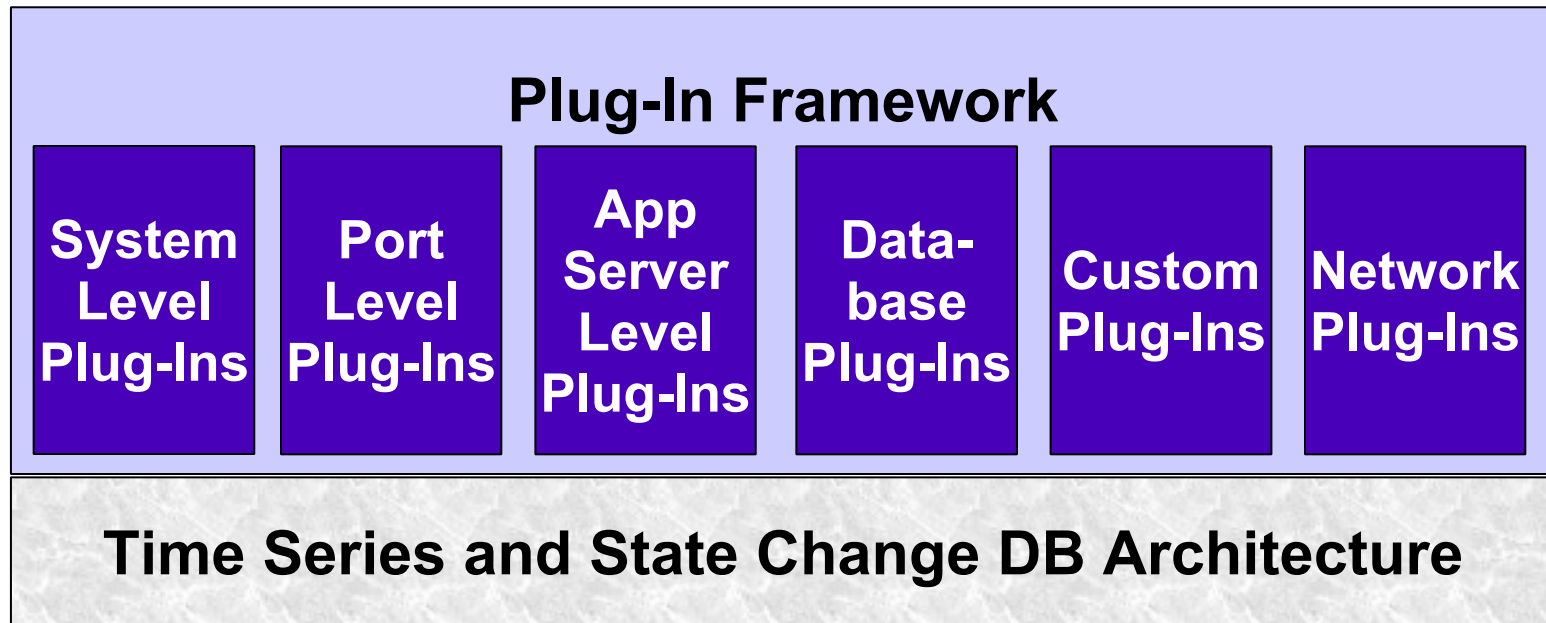
Custom Plug-ins

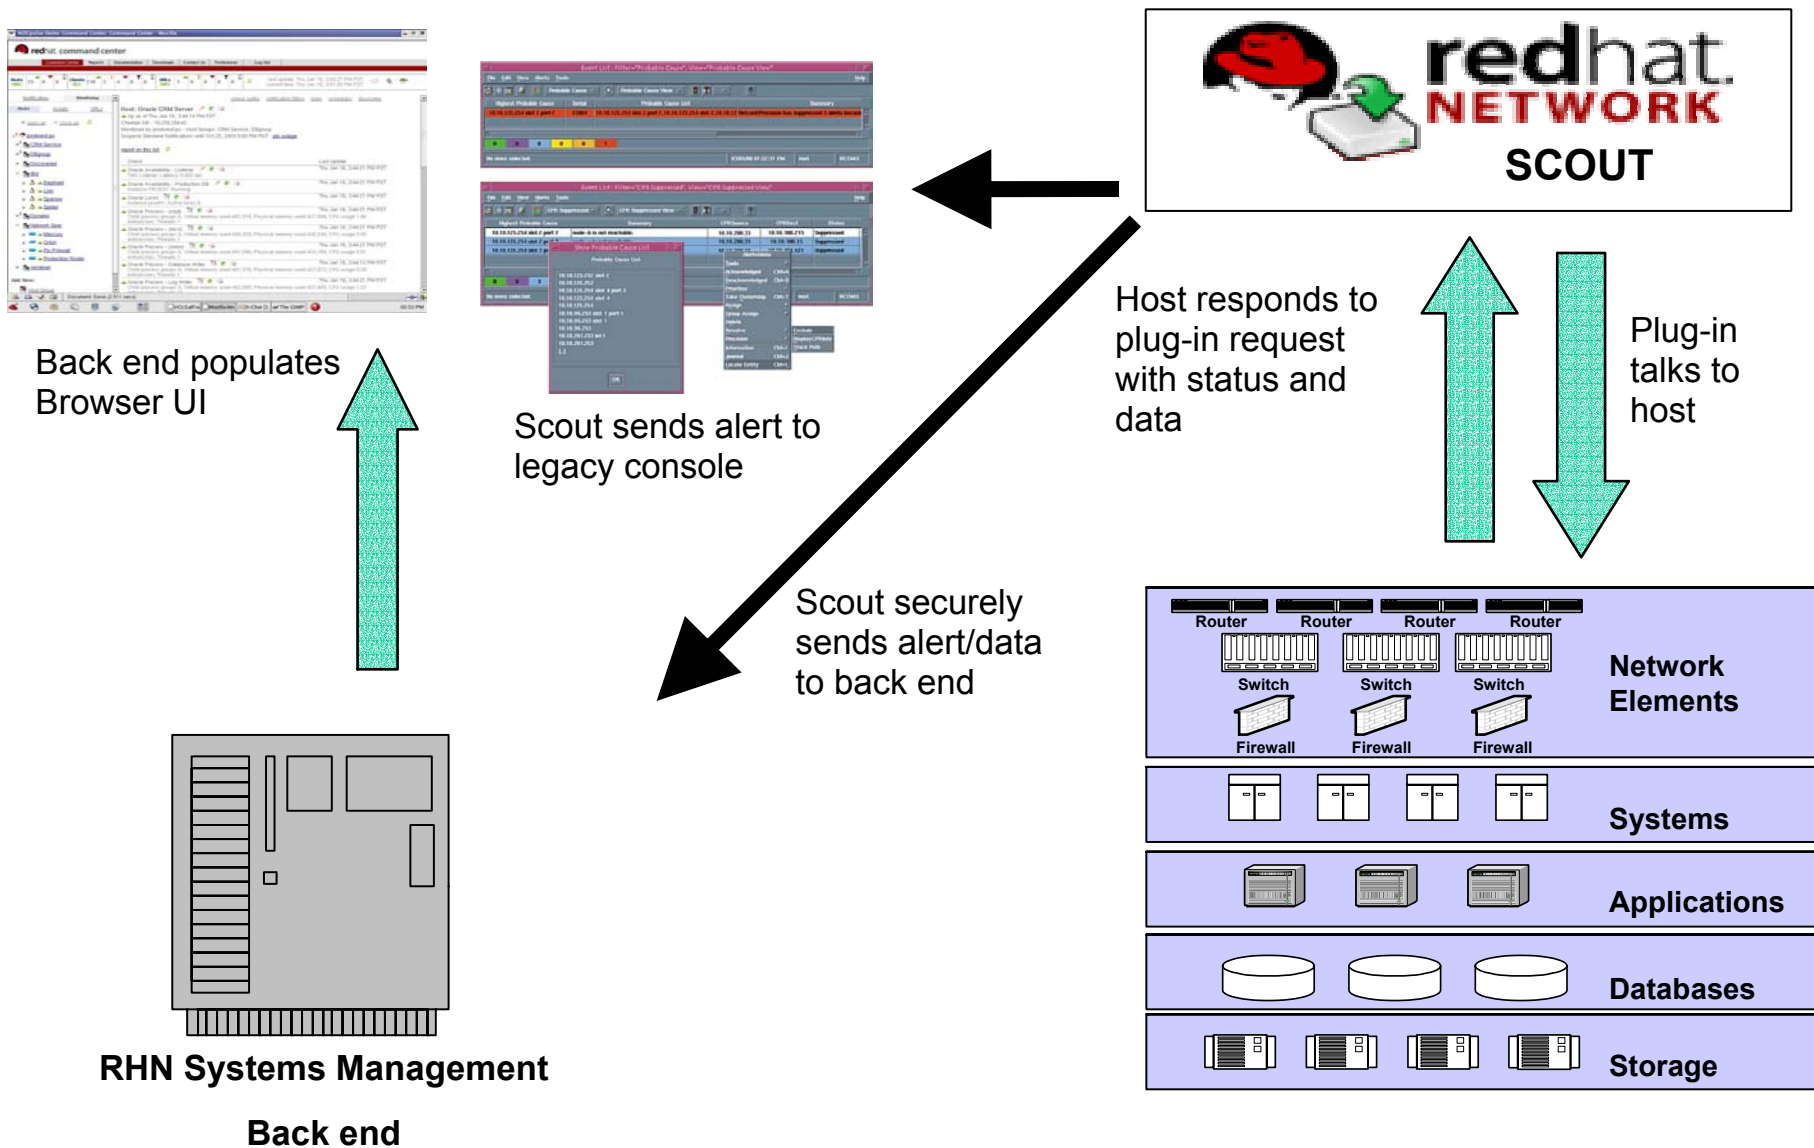
Port, Protocol or
API-based



RHN – Plug-in Framework

Network Plug-ins
Cisco, Alteon/Nortel,
Foundry, F5







Thank You!



Appendix



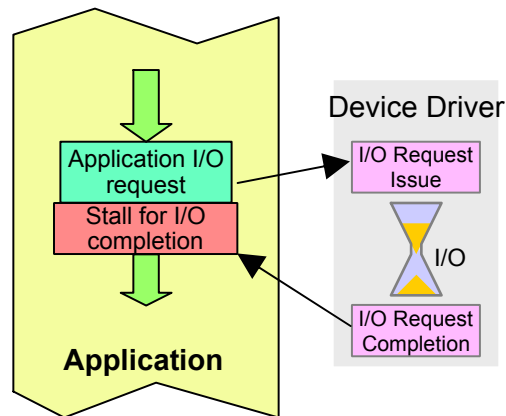
Red Hat Enterprise Linux 2.1 Technical Enhancements



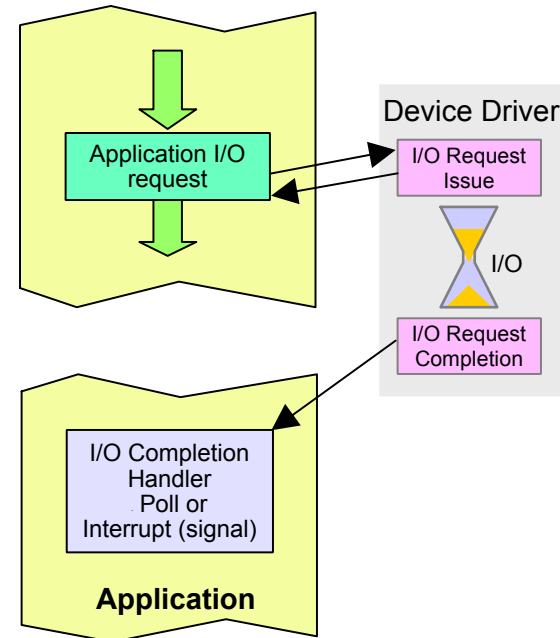
Scalability: Asynchronous I/O

Permits application to continue processing after issuing read I/Os

- Prevents application stalls
- Enables much higher I/O throughput
- Useful for applications moving large of amounts of data
- Supports raw devices, ext2, and ext3



Synchronous I/O

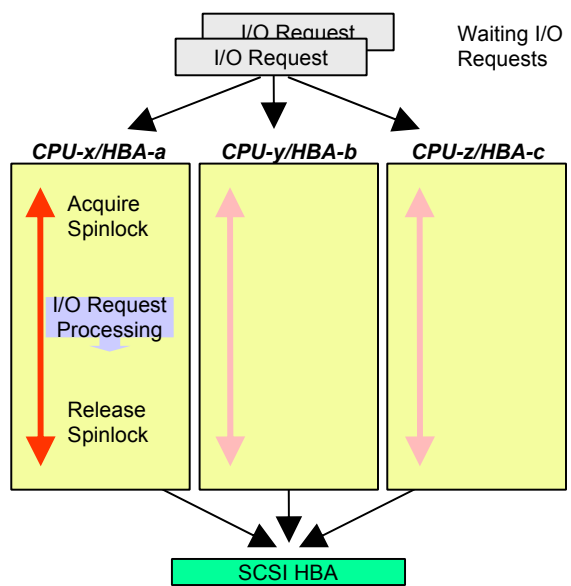


Asynchronous I/O

Scalability: SCSI Adapter Spinlocks

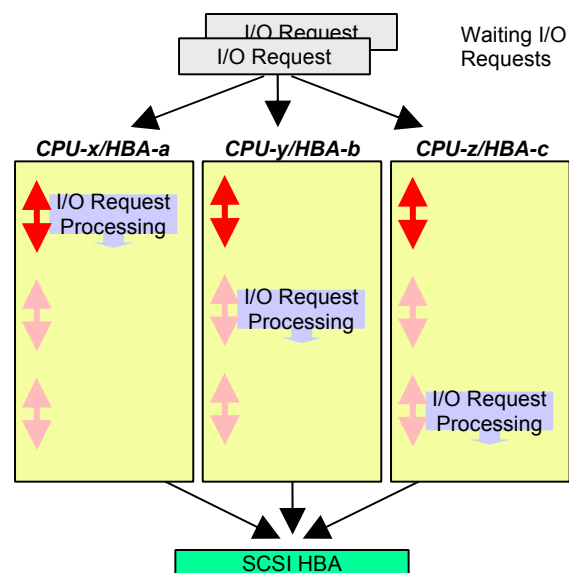
Adds per-SCSI adapter spinlock

- Enables multi-threaded I/O operations in SMP systems
- Support initially provided for Adaptec and Qlogic HBAs



Standard Linux Kernel with coarse SCSI spinlock(s)

SCSI
Device
Driver



RHL-AS enhanced Linux Kernel with fine SCSI spinlock(s)

Scalability: SCSI layer optimization

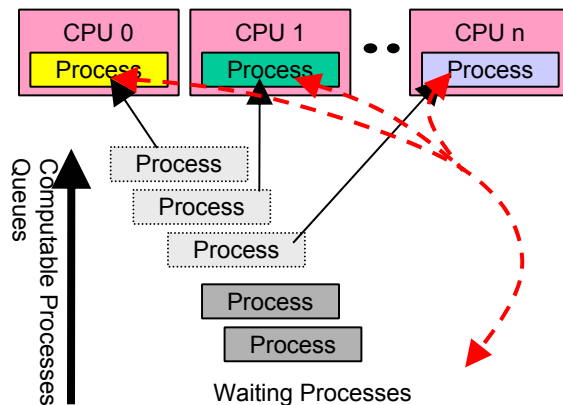
- .SCSI Adapter Spinlocks
 - .Enable multithreaded I/O operations on SMP systems
 - .Reduce contention when using multiple HBAs
- .SCSI driver enhancements
 - .Adaptec aic7xxx (tuning; new locking code)
 - .Adaptec aacraid (rewrite; still using old locking code)
 - .Qlogic qla2200/qla2300 (major cleanup; still using old locking code)
 - .AMI/LSI MegaRAID (major cleanup; new locking code)



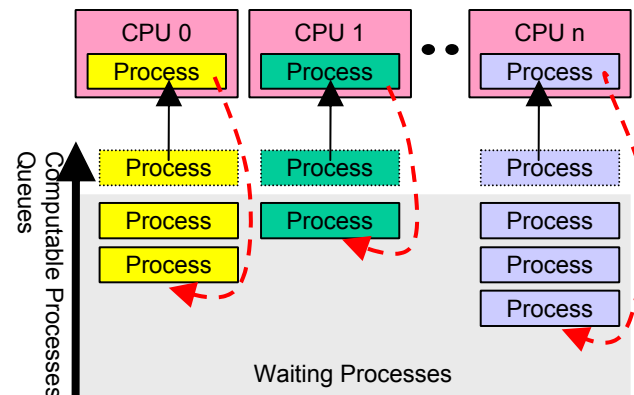
Scalability: $O(1)$ Process Scheduler

Replaces single scheduler compute queue with per-CPU scheduling queue

- Provides processor affinity for SMP systems
- Greatly reduces scheduler spinlock contention compared to current Linux kernels
- Tests show up to 6000% more context switches/second on 8-way SMP system



Standard Linux Kernel with single compute queue



Advanced Server Linux Kernel with multiple compute queues


Bounce Buffer Elimination

Bounce buffers are used by the standard Linux kernel in systems with >1GB of memory

- Bounce (or copy) data from low to high memory

Advanced Server greatly reduces or eliminates bounce operations in systems with >1GB of memory

- Significant performance improvement for I/O intensive applications



Benchmark Results

[Http://ecperf.theserverside.com/ecperf/](http://ecperf.theserverside.com/ecperf/)

- BEA WebLogic 7 and Oracle9i RDBMS
 - **25394.60 BBops/min @ \$9/BBop †**
 - 4 x HP ProLiant ML530-G2 (2 x 2.4 GHz, 2 GB RAM)
 - **16668.47 BBops/min @ \$22/BBop ‡**
 - 1 x Sun Fire V880 (8 x 900 MHz, 16 GB RAM)
- Oracle9iAS and Oracle9i RDBMS
 - **24639.37 BBops/min @ \$5/BBop #**
 - 3 x Compaq ProLiant DL360G2 (2 x 1.4 GHz, 4 GB RAM)
 - **36122.60 BBops/min @ \$12/BBop ‡**
 - 3 x Sun Fire V480 (4 x 900 MHz, 16 GB RAM)



† Red Hat Linux Advanced Server 2.1AS and Red Hat Linux 7.2

‡ Solaris 8

Red Hat Linux Advanced Server 2.1AS

